

IR-168

The complete version (178 pages) of this dissertation is downloadable from the web at:
<http://ciir.cs.uass.edu/publications>

19990719 017

SCALABLE DISTRIBUTED ARCHITECTURES FOR INFORMATION RETRIEVAL

A Dissertation Presented

by

ZHIHONG LU

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

May 1999

Department of Computer Science

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

ACKNOWLEDGMENTS

First and foremost, I would like to express my sincere thanks to my advisor, Kathryn McKinley for her contribution in the development of my research and her guidance in my professional growth. Without her broad vision, deep insight, and strong encouragement, this work would not have been possible.

I would like to thank Bruce Croft and Jamie Callan for teaching me sound research principles on distributed information retrieval and their valuable advice.

I would like to thank the rest of my committee, Don Towsley and Mani Krishna for their valuable comments and suggestions.

I would like to thank all present and past members of the Center for Intelligent Information Retrieval at the University of Massachusetts, Amherst. I was fortunate enough to be a member of this friendly, enjoying, and stimulating research community. In particular, I would like to thank Brendon Cahoon for his preliminary work on distributed information retrieval architectures and his kindness to tutor me the simulation techniques. I would like to thank Jay Ponte for proof-reading my proposal. I would like to thank our secretary, Kate Moruzzi, for her kindness and efficient assistance whenever I needed.

I would like to thank all my friends I have met in Amherst. Without them, the last five and a half years would have been unbearable. In particular, I would like to thank the Chinese community in the Department of Computer Science and the Zinsmeister family.

I would like to thank my husband, Jianmin Wang, for his love, understanding, and absolute faith in my abilities. I owe my son, Luhan Wang, for his four-year life without his mother around. I wish he could understand me when he grows up.

I dedicate this dissertation to my parents, Hao Lu and Jiangqiu Liu, for their love, caring, sacrifice, and support. I will always be grateful to my parents, who have always given me more than what I could ever ask for. I especially appreciate that they teach me the positive attitude to survive in this world. I also dedicate this dissertation to my brother, Qinyi Lu. The dreams we were sharing during our childhood have always been the greatest inspiration for me to move on and face whatever challenge may appear.

This material is based on work supported in part by the National Science Foundation, Library of Congress and Department of Commerce under cooperative agreement number EEC-9209623, and in part by Defense Advanced Research Projects Agency/ITO under ARPA order number D468, issued by ESC/AXS contract number F19628-95-C-0235. Any opinions, findings and conclusions or recommendations expressed in this material are the author and do not necessarily reflect those of the sponsors.

ABSTRACT

SCALABLE DISTRIBUTED ARCHITECTURES FOR INFORMATION RETRIEVAL

MAY 1999

ZHIHONG LU

B.Sc., TONGJI UNIVERSITY

M.Sc., INSTITUTE OF COMPUTING TECHNOLOGY,

CHINESE ACADEMY OF SCIENCES

Ph.D., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Kathryn S. McKinley

As information explodes across the Internet and intranets, information retrieval (IR) systems must cope with the challenge of scale. How to provide scalable performance for rapidly increasing data and workloads is critical in the design of next generation information retrieval systems. This dissertation studies scalable distributed IR architectures that not only provide quick response but also maintain acceptable retrieval accuracy. Our distributed architectures exploit parallelism in information retrieval on a cluster of parallel IR servers using symmetric multiprocessors, and use partial collection replication and selection as well as collection selection to restrict the search to a small percentage of data while maintaining retrieval accuracy.

We first investigate using partial collection replication for IR systems. We examine query locality in real systems, how to select a partial replica based on relevance,

how to load-balance between replicas and the original collection, as well as updating overheads and strategies. Our results show that there exists sufficient query locality to justify partial replication for information retrieval. Our proposed replica selection algorithm effectively selects relevant partial replicas, and is inexpensive to implement. Our evidence also indicates that partial replication achieves better performance than caching queries, because the replica selection algorithm finds similarity between non-identical queries, and thus increases observed locality.

We use a validated simulator to perform a detailed performance evaluation of distributed IR architectures. We explore how best to build parallel IR servers using symmetric multiprocessors, evaluate the performance of partial collection replication and collection selection, and compare the performance of partial collection replication with collection partitioning as well as collection selection. At last we present experiments for searching a terabyte of text. We also examine performance changes when we use fewer large servers, faster servers, and longer queries.

Our results show that because IR systems have heavy computational and I/O loads, the number of CPUs, disks, and threads must be carefully balanced to achieve scalable performance. Our results show that partial collection replication is much more effective at decreasing the query response time than collection partitioning for a loaded system, even with *fewer* resources, and it requires only modest query locality. Our results also show that partial collection replication performs better than collection selection when there exists enough query locality, and it performs worse when the collection access is fairly uniform after collection selection. Finally our results show that replica and collection selection can be combined to provide quick response time for a terabyte of text. Changes of system configurations do not significantly change the relative improvements due to partial collection replication and collection selection, although they affect the absolute response time.